

Fusion trees [Fredman & Willard - JCSS 1993] $O(\log_w n)$ time for predecessor/successor $O(n)$ space - here, static

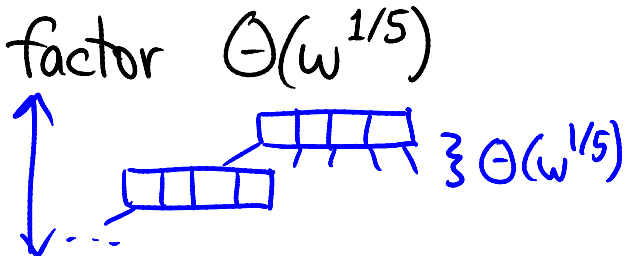
[- dynamic via exponential trees:] [Andersson & Thorup-
 $O(\log_w n + \lg \lg n)$ arXiv:cs.DS/0210006]

Top-level idea:

- B-tree with branching factor $\Theta(w^{1/5})$

\Rightarrow height = $\Theta(\log_w n)$

$$= \Theta\left(\frac{\lg n}{\lg w}\right)$$



- search must visit each node in $O(1)$ time

- need to find correct branch

- not enough time to read node ($\Theta(w^{1/5})$ words)

Fusion-tree node:

given k keys $x_0 < x_1 < \dots < x_{k-1}$, $k = O(w^{1/5})$

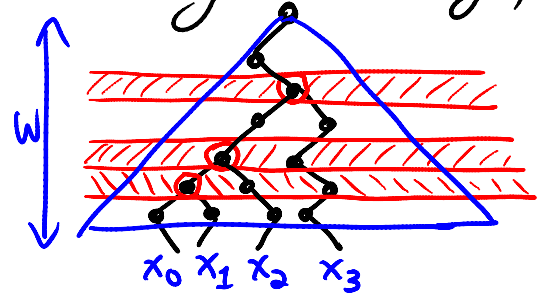
preprocess in $k^{O(1)}$ time

subject to predecessor/successor in $O(1)$ time

Note: $\min \left\{ \underbrace{\log_w n}_{\text{fusion}}, \underbrace{\lg w}_{\text{van Emde Boas}} \right\} \leq \sqrt{\lg n}$

Distinguishing $k = O(w^{1/5})$ keys:

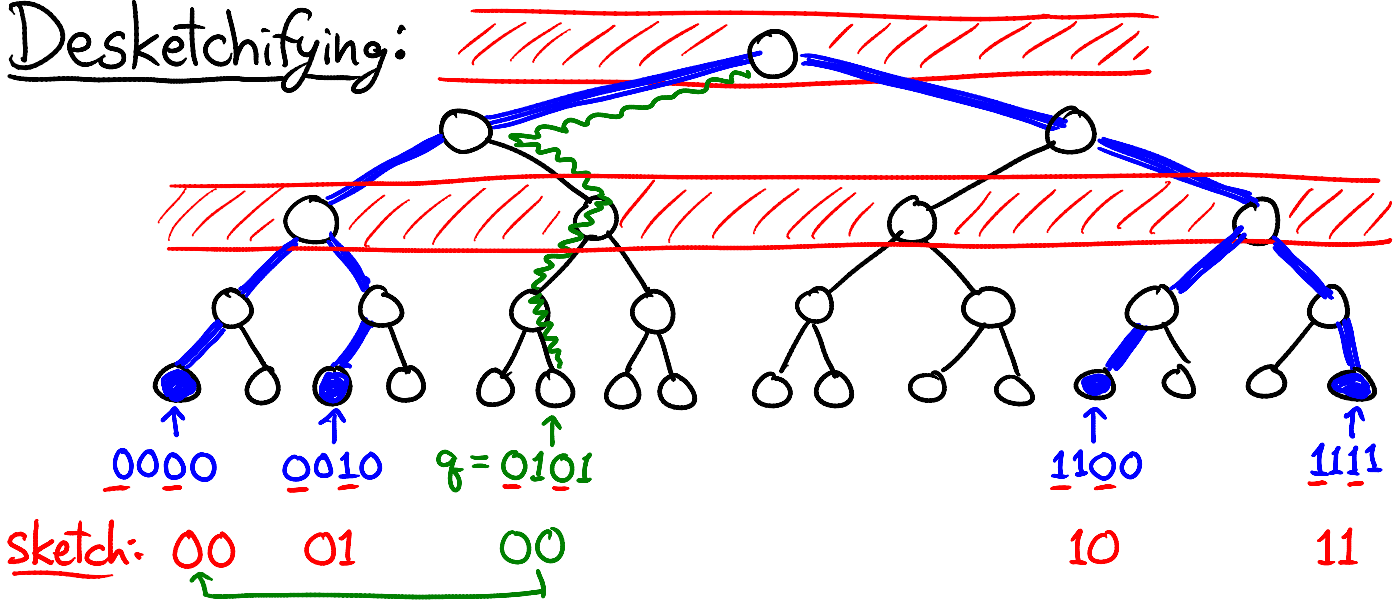
- view keys x_0, x_1, \dots, x_{k-1} as binary strings (0/1)
i.e. root-to-leaf paths in height- w binary tree (left/right)
- \Rightarrow $k-1$ branching nodes \odot
- $\Rightarrow \leq k-1$ levels containing branching nodes
- = bits where x_0, x_1, \dots, x_{k-1}
first differ (first distinct prefix)
- call these important bits $b_0 < b_1 < \dots < b_{r-1}$, $r < k = O(w^{1/5})$



(perfect-)sketch(x) = extract bits b_0, b_1, \dots, b_r from word x
i.e. r -bit vector whose i th entry = b_i th bit of x
 \Rightarrow $\text{sketch}(x_0) < \dots < \text{sketch}(x_k)$ ~ all different & same order
& all packable into a word: $k \cdot r = O(w^{2/5})$ bits
- computable in $O(1)$ time as an AC^0 operation
[Andersson, Miltersen, Thorup - TCS 1999]

Idea: for query q , compare $\text{sketch}(q)$ in parallel
to $\text{sketch}(x_0), \dots, \text{sketch}(x_{k-1})$
- again AC^0 operation on $O(1)$ words
 \Rightarrow find where $\text{sketch}(q)$ fits (predecessor/successor)
among $\text{sketch}(x_0) < \dots < \text{sketch}(x_{k-1})$

Desketchifying:

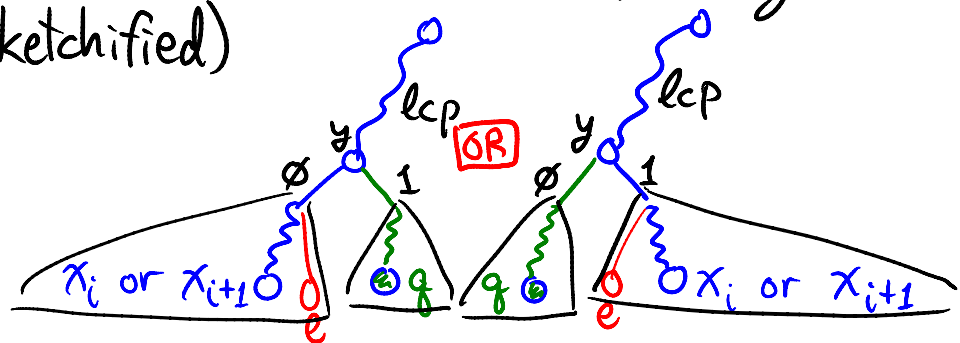


- suppose $sketch(x_i) \leq sketch(q) < sketch(x_{i+1})$
- compute longest common prefix = lowest common ancestor between q and $(x_i \text{ or } x_{i+1})$ — take longest/lowest

= node y where q fell off the "right paths" to x_i 's
 \Rightarrow correct direction/subtree = $\begin{cases} y1 & \text{if } (y+1)\text{st bit of } q = 0 \\ y0 & \text{if } (y+1)\text{st bit of } q = 1 \end{cases}$

- extreme e in subtree closest to $q = \begin{cases} y100\dots0 \\ y011\dots1 \end{cases}$

- predecessor & successor of q among x_i 's
- = predecessor & successor of $sketch(e)$ among $sketch(x_i)$'s (desketchified)



Parallel comparison:

- $\text{sketch}(\text{node}) = 1 \text{ sketch}(x_0) 1 \text{ sketch}(x_1) \dots 1 \text{ sketch}(x_{k-1})$
- $\text{sketch}(q)^k = 0 \text{ sketch}(q) 0 \text{ sketch}(q) \dots 0 \text{ sketch}(q)$
 $= \text{sketch}(q) \cdot 000001 \ 000001 \ \dots \ 000001$
- difference $= \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{*****} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{*****} \ \dots \ \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{*****}$
- AND with $100000 \ 100000 \ \dots \ 100000$
 $\rightarrow \begin{pmatrix} 1 \\ 0 \end{pmatrix} 00000 \ \begin{pmatrix} 1 \\ 0 \end{pmatrix} 00000 \ \dots \ \begin{pmatrix} 1 \\ 0 \end{pmatrix} 00000$

Idea: extra 1's in $\text{sketch}(\text{node})$ protect from underflow

- remains 1 if $\text{sketch}(q) \leq \text{sketch}(x_i)$
- becomes 0 if $\text{sketch}(q) > \text{sketch}(x_i)$

$\Rightarrow \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ bits are 0 0 0 0 1 1 1 1
where $\text{sketch}(q)$ fits $\uparrow \leftarrow$ most significant 1 bit

Index of most significant 1 bit: 00010110 \mapsto 4
7 6 5 4 3 2 1 0

- AC⁰ operation [Andersson, Miltersen, Thorup - TCS 1999]
- instruction on many real CPUs e.g. Pentiums
(see Linux kernel: `include/asm-*/bitops.h`)
- ^{LSB} possible with $O(1)$ arithmetic & Boolean ops. [Brodnik 1993]
- easy solution with extra space:
 - lookup table on all strings of $\epsilon \cdot w$ bits
 $\Rightarrow O(2^{\epsilon \cdot w} \lg w)$ bits = $O(u^\epsilon \lg \lg u)$ bits of space
 - query: $O(1/\epsilon)$ [or $O(\lg(1/\epsilon))$] probes to table
- note for next time: cover Fredman & Willard's clever solution

Approximate sketch(x) \Rightarrow word RAM

- don't need sketch to pack b_i bits right next to each other
- can be spread out in predictable pattern of length $O(w^{4/5})$
independent of x

Idea: mask important bits: $x' = x \text{ AND } \sum_{i=0}^{r-1} 2^{b_i}$
& multiply $x' \cdot m = \left(\sum_{i=0}^{r-1} x_{b_i} 2^{b_i} \right) \cdot \left(\sum_{j=0}^{r-1} 2^{m_j} \right)$
$$= \sum_{i=0}^{r-1} \sum_{j=0}^{r-1} x_{b_i} 2^{b_i + m_j}$$

Claim: for any b_0, b_1, \dots, b_{r-1} , can choose m_0, m_1, \dots, m_{r-1} such that

- (a) $b_i + m_j$ are all distinct \Rightarrow no collision
- (b) $b_0 + m_0 < b_1 + m_1 < \dots < b_{r-1} + m_{r-1} \Rightarrow$ preserve order
- (c) $(b_{r-1} + m_{r-1}) - (b_0 + m_0) = O(r^4) = O(w^{4/5}) \Rightarrow$ small span

\Rightarrow approx-sketch(x) = $\left((x \cdot m) \text{ AND } \underbrace{\sum_{i=0}^{r-1} 2^{b_i + m_i}}_{\text{discard } i \neq j} \right) \gg (b_0 + m_0)$

Proof: ① choose $m'_0, m'_1, \dots, m'_{r-1} < r^3$ such that $b_i + m'_j$ are all distinct modulo r^3 (strong (a))

- pick $m'_0, m'_1, \dots, m'_{t-1}$ by induction
- m'_t must avoid $\underbrace{m'_i}_{t} + \underbrace{b_j}_{r} - \underbrace{b_k}_{r} \forall i, j, k$
 $\Rightarrow tr^2 < r^3$ choices

\Rightarrow choice for m'_t exists

② let $m_i = m'_i + (w - b_i + i r^3)$ rounded down to multiple of r^3
 $\equiv m'_i \pmod{r^3}$

$\Rightarrow m_i + b_i$ in r^3 interval after $(\frac{w}{r^2} + i)$ th multiple of r^3

$\Rightarrow \underbrace{m_0 + b_0}_{\approx w} < m_1 + b_1 < \dots < \underbrace{m_{r-1} + b_{r-1}}_{\approx w + r^4} \Rightarrow$ difference = $O(r^4)$